arXiv:1405.5645v1 [cs.LO] 22 May 2014

A Variant of Earley Deduction With Partial Evaluation

Heike Stephan and Stefan Brass

Martin-Luther-Universität Halle-Wittenberg, Institut für Informatik,
Von-Seckendorff-Platz 1, D-06099 Halle (Saale), Germany
stephan@informatik.uni-halle.de, brass@informatik.uni-halle.de

Abstract. We present an algorithm for query evaluation given a logic program consisting of function-free Datalog rules. It is based on Earley Deduction [4,6] and uses a partial evaluation similar to the one we developed for our SLDMagic method [1]. With this, finite automata modeling the evaluation of given queries are generated. In certain cases, the new method is more efficient than SLDMagic and the standard Magic Set method since it can process several deduction steps as one.

1 Introduction

The goal of deductive database systems is to offer integrated systems which permit to do programming and database tasks in a single, declarative language. This would improve the current situation in which several languages are mixed, e. g. Java and SQL. While SQL is declarative and has successfully shown the advantages of declarative languages, the language used for application programming is usually non-declarative.

Whereas in earlier times deductive database research was concentrated only on recursive query evaluation, now new applications, e. g. for the semantic web, are in the focus. But even so mundane tasks as the generation of web pages must be considered if deductive databases should be used for real world application programming. In [3] we made a proposal for declarative output and also investigated how sorting can be integrated in Datalog, which is obviously important for output and also for database queries.

However, even the classical task of bottom-up query evaluation deserves more research in order to improve efficiency [2]. The new method presented in this paper loops through a sequence of states (sets of rules being processed). From one state to the next, the successor state is determined by a single database fact and the preceding rule set. At compilation time, when the database state is not yet known, partial evaluation of the program is done by using facts with abstract values. The method is based on Earley Deduction [4,6] which exploits the similarity of context free grammar rules and rules of logic programs. The partial evaluation is similar to the one we developed for our SLDMagic method [1]. It makes the algorithm very competetive because program analysis and abstract execution can already be done at compilation time.

The algorithm is also interesting because it especially fits applications in which input must be parsed (after all, the Earley algorithm is a parsing algorithm). An input text, e. g. `"abc"`, can be represented by Datalog facts as follows:

```
input(1, a, 2).
input(2, b, 3).
input(3, c, 4).
eof(4).
```

This is very similar to the standard difference list technique for definite clause grammars, but since Datalog has no lists we use positions in the input.

First we give some basic definitions (Section 2), then we describe the basic deducion method (Section 3) and finally present the partial evaluation (Section 4).

2 Basic Definitions

Definition 1 (Rule). *A rule is a formula of the form* $\mathsf{A} \leftarrow \mathsf{B}_1 \wedge \cdots \wedge \mathsf{B}_n$ *where* A *and* $\mathsf{B}_i, i = 1, \ldots, n$, *are positive literals, i. e. atomic formulas* $\mathsf{p}(t_1, \ldots, t_m)$ *with a predicate* p *of arity* m *and terms* $t_j, j = 1, \ldots, m$. *Terms are variables or constants. In the above rule,* A *is called the head and* $\mathsf{B}_1 \wedge \cdots \wedge \mathsf{B}_n$ *is called the body. A rule with empty body (i. e.* $n = 0$*) and without variables is called a fact.*

In the context of deductive databases, the *range restriction* condition ensures that no derived fact contains variables. In the remaining paper, this condition is assumed to be satisfied for every rule.

Definition 2 (Range Restriction). *A rule is range restricted iff every variable that appears in the head appears also in the body.*

Definition 3 (EDB- and IDB-Predicates, Program, and Database). *Predicates are partitioned into EDB ("extensional database") predicates defined by facts and IDB ("intensional database") predicates defined by rules. A logic program is a finite set of rules with an IDB predicate in the head and at least one body literal. A database is a finite set of facts with EDB predicate.*

The requirement that the body of a program rule is non-empty simplifies later definitions but is no restriction: One can use a special EDB predicate true without arguments in the rule body.

Definition 4 (Answer Predicate, Goal Rule, and Query). *We assume that an IDB predicate* answer *is distinguished as "main" predicate. It must not appear in the body of a program rule. A rule with the predicate* answer *in the head is called goal rule and represents a query to the logic program.*

The goal of query evaluation is to determine the answer-facts which are derivable from program and database together.

3 Deduction Method

The deduction method uses sequences of states to compute facts of the answer relation.

Definition 5 (Rule Normalization). *Let VAR be the set of variables of a rule* R*, and let this set be ordered by the occurrence of its elements in* R*:* $VAR = \{V_i \mid i \in \{0, \ldots, |VAR|-1\}\}$ *where* $V_j < V_k$ *iff the first occurrence of* V_j *is before the first occurrence of* V_k *and* $j < k$ *iff* $V_j < V_k$ *(*$j, k \in \{0, \ldots, |VAR|-1\}$*). Let further* $\mathcal{X} = \{\mathsf{X}_i \mid i \in \mathbb{N} \cup \{0\}\}$ *be an ordered set of variables:* $\mathsf{X}_j < \mathsf{X}_k$ *iff* $j < k$ *(*$j, k \in \{\mathbb{N} \cup \{0\}$*). Then rule R is normalized by substituting every* $V_i \in VAR$ *by* $\mathsf{X}_i \in \mathcal{X}$ *(*$i \in \{0, \ldots, |VAR|-1\}$*).*

Definition 6 (State). *A state is a set of normalized rules.*

Definition 7 (Selection Function). *A selection function chooses for every rule* $\mathsf{A} \leftarrow \mathsf{B}_1, \ldots, \mathsf{B}_n$ *with* $n \geq 1$ *an index* $i \in \{1, \ldots, n\}$ *(i. e. a body literal).*

In every rule of a state one body literal is selected; for simplicity of presentation it is assumed that this is the leftmost body literal. However, we note that in the database context the selection function is an important optimization parameter. So a real implementation will use a selection function that tries to make use of input constants and possibly existing indexes or other database access structures.

From the rules in a state new rules are derived by two basic derivation steps that are already described in [4], a "downward" *instantiation* and an "upward" *reduction* step. In a way, this can be viewed as splitting up the SLD-resolution step, which avoids deriving rules with arbitrary length. A derived rule is first normalized before added to a state.

If the selected literal of a rule in the state unifies with the head of a program rule, an *instance* of the program rule is derived by renaming all variables in the program rule and applying the most general unifier of the selected literal and the program rule head to the program rule. Thus, instantiation corresponds to calling an IDB predicate as in Prolog's four port box model. Several instances can be derived from the same rule.

Definition 8 (Instance, Instantiation). *Let* $\mathsf{R} = \mathsf{A} \leftarrow \mathsf{B}_1 \wedge \cdots \wedge \mathsf{B}_n$ *be a rule of the program and* $\mathsf{K} \leftarrow \mathsf{L}_1 \wedge \cdots \wedge \mathsf{L}_m, m > 0$*, a rule in the state with selected literal* L_1*. Let* R' *be the rule resulting from* R *by a renaming of variables so that no variable in* R' *occurs in a rule in the state, i. e. there is a substitution* θ *so that* $\mathsf{R}' = \mathsf{R}\,\theta = \mathsf{A}' \leftarrow \mathsf{B}'_1 \wedge \cdots \wedge \mathsf{B}'_n$*. A rule* R'' *is an instance of* R *iff* L_1 *and* A' *are unifiable with most general unifier* σ *and* $\mathsf{R}'' = \mathsf{R}'\,\sigma$*.*

A reduction is performed with a fact, either of the database or a derived one. If there is a derived rule with a selected literal that unifies with the fact, this rule is *reduced* by the fact and a new rule, the *reduct*, is created by applying the most general unifier and removing the selected literal. When the last body literal is removed by reduction, an IDB fact results. Thus, reduction is a special case of resolution with a fact. Again, one fact can be used for several reductions.

Definition 9 (Reduct, Reduction). *Given a derived rule* $\mathsf{R} = \mathsf{A} \leftarrow \mathsf{B}_1 \wedge \mathsf{B}_2 \wedge \cdots \wedge \mathsf{B}_n$ *with selected literal* B_1 *and a fact* F *in the database or in the state, the rule* R' *is a reduct of* R *iff* B_1 *and* F *are unifiable with most general unifier* σ *such that* $\mathsf{R}' = (\mathsf{A} \leftarrow \mathsf{B}_2 \wedge \cdots \wedge \mathsf{B}_n)\,\sigma$*. The corresponding derivation step is called reduction,* F *reduces* R *to* R'*.*

Definition 10 (Initial State). *The initial state consists of the goal rule and all rules that can be iteratively derived by instantiation.*

At a state transition, exactly one EDB fact is used to compute the successor state.

Definition 11 (Dependency-Relation of Rules). *A rule* R *depends directly on a rule* R' *iff the selected literal in* R *is unifiable with the head literal of* R'*. A rule* R *depends on a rule* R' *with respect to a state* $\mathcal{S}$ *iff there are rules* $\mathsf{R}_1, \ldots, \mathsf{R}_n \in \mathcal{S}$ *such that* $\mathsf{R}_1 = \mathsf{R}$*, each* $\mathsf{R}_i, i = 1, \ldots, n-1$*, depends directly on* R_{i+1}*, and* R_n *depends directly on* R' *(note that* R' *does not have to be contained in* $\mathcal{S}$*).*

Definition 12 (Successor State). *Let a program* $\mathcal{P}$*, a database* $\mathcal{D}$*, a state* $\mathcal{S}$*, and a fact* $\mathsf{F} \in \mathcal{D}$ *be given. The successor state* $\mathcal{S}'$ *is constructed as follows:*

1. *First,* $\mathcal{S}'$ *is initialized with all rules that result from reduction applied to rules in* $\mathcal{S}$ *with fact* F*. If the result is empty, there is no successor state.*
2. *If* $\mathcal{S}'$ *now contains IDB facts, reduction is applied repeatedly to rules in* $\mathcal{S}$ *with facts in* $\mathcal{S}'$ *and the results are inserted into* $\mathcal{S}'$ *until nothing changes.*
3. *Then instantiation is applied iteratively to each rule* $\mathsf{R} \in \mathcal{S}'$ *with a selected IDB-literal. All instances are added to the successor state.*
4. *Finally, rules* $\mathsf{R} \in \mathcal{S}$ *that depend (with respect to* $\mathcal{S}$*) on a rule with at least one body literal in* $\mathcal{S}'$ *are copied to* $\mathcal{S}'$*. The copied rules are those that still have a chance of being reduced by an IDB fact.*

Definition 13 (State Sequence). *States* $\mathcal{S}_1, \ldots, \mathcal{S}_n$ *form a state sequence iff every* $\mathcal{S}_{i+1}$ *is the successor state for* $\mathcal{S}_i$ *and a fact* F_i *of the database,* $i = 1, \ldots, n-1$*.*

Definition 14 (Computed Answers). *A fact* $\mathsf{answer}(c_1, \ldots, c_m)$ *is computed if there is a state sequence* $\mathcal{S}_1, \ldots, \mathcal{S}_n$ *such that* $\mathcal{S}_1$ *is the initial state and* $\mathsf{answer}(c_1, \ldots, c_m) \in \mathcal{S}_n$*.*

There can only be finitely many different states for a given program $\mathcal{P}$ and database $\mathcal{D}$ for the following reasons:

- The state contains only predicates and constants occurring in the finite set $\mathcal{P} \cup \mathcal{D}$.
- No derived rule can become longer than the longest program rule.
- A state does not contain two rules that differ only in the names of their variables.

However, the state sequence could be cyclic, so one must check whether a newly constructed state is indeed new. Of course, optimizations are possible and subject of our further research.

Example 1. Let the left recursive version of the standard transitive closure program be given:

[1] $\mathsf{path}(\mathsf{X}_0, \mathsf{X}_1) \leftarrow \mathsf{edge}(\mathsf{X}_0, \mathsf{X}_1)$.
[2] $\mathsf{path}(\mathsf{X}_0, \mathsf{X}_1) \leftarrow \mathsf{path}(\mathsf{X}_0, \mathsf{X}_2) \wedge \mathsf{edge}(\mathsf{X}_2, \mathsf{X}_1)$.

Let the database be

[3] $\mathsf{edge}(1, 2)$.
[4] $\mathsf{edge}(2, 3)$.

Now let the following goal rule be given:

[5] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_0)$.

The initial state $\mathcal{S}_0$ consists of the goal rule plus rules added by instantiation:

[6] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_0)$. //goal [5]
[7] $\mathsf{path}(1, \mathsf{X}_0) \leftarrow \mathsf{edge}(1, \mathsf{X}_0)$. //inst. of [1] because of [6]
[8] $\mathsf{path}(1, \mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_1) \wedge \mathsf{edge}(\mathsf{X}_1, \mathsf{X}_0)$. //inst. of [2] because of [6]

Rule [8] also calls for instantiation but that gives again [7] and [8].

Now there is only one database fact, $\mathsf{edge}(1, 2)$, that leads to a successor state, and by reducing with this fact we reach state $\mathcal{S}_1$:

[9] $\mathsf{path}(1, 2)$. //Reduction of [7] with [3]
[10] $\mathsf{answer}(2)$. //Reduction of [6] with [9]
[11] $\mathsf{path}(1, \mathsf{X}_0) \leftarrow \mathsf{edge}(2, \mathsf{X}_0)$. //Reduction of [8] with [9]
[12] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_0)$. //Copy of [6] because of [11]
[13] $\mathsf{path}(1, \mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_1) \wedge \mathsf{edge}(\mathsf{X}_1, \mathsf{X}_0)$. //Copy of [8] because of [11]

Again, reduction with only one database fact, $\mathsf{edge}(2, 3)$, is possible and gives the state $\mathcal{S}_2$:

[14] $\mathsf{path}(1, 3)$. //Reduction of [11] with [4]
[15] $\mathsf{answer}(3)$. //Reduction of [12] with [14]
[16] $\mathsf{path}(1, \mathsf{X}_0) \leftarrow \mathsf{edge}(3, \mathsf{X}_0)$. //Reduction of [13] with [14]
[17] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_0)$. //Copy of [12] because of [16]
[18] $\mathsf{path}(1, \mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_1) \wedge \mathsf{edge}(\mathsf{X}_1, \mathsf{X}_0)$. //Copy of [13] because of [16]

No more reductions with database facts can be applied to rules in $\mathcal{S}_2$. □

Theorem 1 (Correctness). *Let a program $\mathcal{P}$ and a database $\mathcal{D}$ be given. Every computed answer is indeed a logical consequence of $\mathcal{P} \cup \mathcal{D}$.*

Proof. This is easy: Each step (reduction and instantiation) is a logical consequence of $\mathcal{P} \cup \mathcal{D}$ and the previously computed rules. □

Theorem 2 (Completeness). *For every ground substitution θ such that* $\mathsf{answer}(\mathsf{X}_1, \ldots, \mathsf{X}_q)\,\theta$ *is a logical consequence of the program and the database,* $\mathsf{answer}(\mathsf{X}_1, \ldots, \mathsf{X}_q)\,\theta$ *is computed.*

The completeness theorem is a corollary of the following lemma, if $\mathcal{S}_0$ is the initial state and the rule considered is the goal rule.

Lemma 1. *Let a program $\mathcal{P}$ and a database $\mathcal{D}$ be given. If a state $\mathcal{S}_0$ contains a rule* $\mathsf{R} = \mathsf{A} \leftarrow \mathsf{B}_1 \wedge \cdots \wedge \mathsf{B}_n$ *and there is a ground substitution θ such that each* $\mathsf{B}_i\,\theta$ *is a logical consequence of $\mathcal{P} \cup \mathcal{D}$, then there is a state sequence $\mathcal{S}_0, \mathcal{S}_1, \ldots, \mathcal{S}_m$ such that* $\mathsf{A}\,\theta$ *is contained in $\mathcal{S}_m$. Furthermore, any rule* $\mathsf{R}' \in \mathcal{S}_0$ *that depends on* R *is contained in every state $\mathcal{S}_1, \ldots, \mathcal{S}_{m-1}$.*

Proof. Since the $\mathsf{B}_i\,\theta$ are logical consequences of $\mathcal{P} \cup \mathcal{D}$, they are contained in the least fixpoint of $\mathsf{T}_{\mathcal{P}\cup\mathcal{D}}$, and because there are no function symbols, this is reached after a finite number of iterations. The proof is by induction on the maximum (over i, $i = 1, \ldots, n$) of the number of steps needed to derive $\mathsf{B}_i\,\theta$ with the $\mathsf{T}_{\mathcal{P}\cup\mathcal{D}}$-operator.

If this is 1, all $\mathsf{B}_i\,\theta$ are contained in $\mathcal{D}$. For proving the first step of induction, there is a second induction on n (the number of body literals in rule R). If this is 1, then $\mathsf{R} = \mathsf{A} \leftarrow \mathsf{B}_1$, $\mathcal{S}_m = \mathcal{S}_1$ is the direct successor of $\mathcal{S}_0$ and contains $\mathsf{A}\,\theta$. $\mathcal{S}_0$ contains all rules depending on R. Now assume that the theorem is proven for rules R with n body literals B_i, all of them with EDB predicate. Assume further that $\mathcal{S}_0$ contains $\mathsf{A} \leftarrow \mathsf{B}_1 \wedge \mathsf{B}_2 \wedge \cdots \wedge \mathsf{B}_{n+1}$. If B_1 is selected, there exists a direct successor state of $\mathcal{S}_0$ with the fact $\mathsf{B}_1\,\theta_1$ that contains $(\mathsf{A} \leftarrow \mathsf{B}_2 \wedge \cdots \wedge \mathsf{B}_{n+1})\,\theta_1$ and all rules in $\mathcal{S}_0$ depending on this rule, where θ_1 is θ restricted to the variables occurring in B_1. From the hypothesis of the second induction the theorem follows, and $m = n + 1$.

Now assume that the theorem is proven for all cases where the body literals $\mathsf{B}_i\,\theta$ are derivable after at most k steps of the $\mathsf{T}_{\mathcal{P}\cup\mathcal{D}}$-operator. This means that $\mathsf{A}\,\theta$ can be derived after $k+1$ steps and computed with a state sequence $\mathcal{S}_0, \ldots, \mathcal{S}_m$ of length $m + 1$.

For the induction step, suppose that all $\mathsf{B}_i\,\theta$ are derivable after at most $k+1$ steps of the $\mathsf{T}_{\mathcal{P}\cup\mathcal{D}}$-operator. Again, to prove the theorem there is an induction on the number of B_i in R. If this is 1, then $\mathsf{R} = \mathsf{A} \leftarrow \mathsf{B}_1$. We consider only the case that B_1 is an IDB literal (the other case is already shown above). In this case, an instantiation is performed, so $\mathcal{S}_0$ contains a rule $\mathsf{B} \leftarrow \mathsf{C}_1 \wedge \cdots \wedge \mathsf{C}_l$ where B unifies with B_1. This rule is either the new instance or a rule already present in the state and equal to the new instance. From the inductive hypothesis follows that $\mathsf{B}\,\theta_1$ (where θ_1 is θ restricted to the variables occurring in B) can be computed with a state sequence $\mathcal{S}_0, \ldots, \mathcal{S}_m$, and that $\mathcal{S}_{m-1}$ contains the rule $\mathsf{R} = \mathsf{A} \leftarrow \mathsf{B}_1$ which depends on B. Since B_1 is the only body literal, $\theta_1 = \theta$. Thus, in state $\mathcal{S}_m$ a reduction with $\mathsf{B}\,\theta$ and $\mathsf{A} \leftarrow \mathsf{B}_1 \in \mathcal{S}_{m-1}$ can be performed so that $\mathcal{S}_m$ also contains $\mathsf{A}\,\theta$.

If $\mathcal{S}_0$ contains a rule $\mathsf{A} \leftarrow \mathsf{B}_1 \wedge \mathsf{B}_2 \wedge \cdots \wedge \mathsf{B}_n$ where B_1 is selected and an IDB literal, again it also contains a rule $\mathsf{B} \leftarrow \mathsf{C}_1 \wedge \cdots \wedge \mathsf{C}_l$. A state sequence $\mathcal{S}_0, \ldots, \mathcal{S}_{m'}$ can be computed where $\mathcal{S}_{m'}$ contains $\mathsf{B}\,\theta_1$, $\mathcal{S}_{m'-1}$ contains $\mathsf{A} \leftarrow \mathsf{B}_1 \wedge \mathsf{B}_2 \wedge \cdots \wedge \mathsf{B}_n$

and, after a reduction with the fact $\mathsf{B}\,\theta_1$, $\mathcal{S}_{m'}$ contains $\mathsf{A} \leftarrow \mathsf{B}_2 \wedge \cdots \wedge \mathsf{B}_n\,\theta_1$ as well as all rules in $\mathcal{S}_{m'-1}$ depending on this rule. Finally, from both inductive hypotheses the theorem follows. □

4 Partial Evaluation

Especially in database context, the facts of the extensional database might not be known before execution time, and as the aim is to compile a program beforehand, an abstraction from actual data values must be developed. For this purpose, abstract values taken from an infinite set of symbolic constant values, $\mathcal{V}$, are used instead of the data values that are known only at execution time. Constants in program rules and in the query are not substituted, so no symbolic value may occur in a program rule or in the query. Via a partial evaluation an automaton can be constructed that models the process of query evaluation.

We only need to redefine a state transition. For a given state, create a set with all selected EDB literals in the state that are not equal to each other. Two literals are considered equal if they differ only in the names of their variables. For every literal in this set there is a state transition assigned to it which is labeled with the literal. Choose a literal and substitute its variables with new symbolic values that have not yet been used elsewhere. The resulting "symbolic fact" represents all facts that could be obtained from a query to the corresponding EDB relation at execution time. Thus, a state transition can be viewed as a data retrieving interface. Now reduce all possible rules in the state with the symbolic fact and add the derived rules to a new state. If a state contains a fact of the answer relation, it is a *final state*. In the same way as an EDB literal with symbolic values can be viewed as a representative of a set of facts, a state with symbolic values can be viewed as a representative of a set of states that depends on the actual data values.

Example 2. Consider again the transitive closure program with the same goal rule. The database is irrelevant now, only name and arity of EDB predicates are needed. Let the set $\mathcal{V}$ of symbolic values be $\{\mathsf{c}_0, \mathsf{c}_1, \ldots\}$. For the initial state $\mathcal{S}_0$ there are no differences.

[6] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_0)$.
[7] $\mathsf{path}(1, \mathsf{X}_0) \leftarrow \mathsf{edge}(1, \mathsf{X}_0)$.
[8] $\mathsf{path}(1, \mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_1) \wedge \mathsf{edge}(\mathsf{X}_1, \mathsf{X}_0)$.

In rule [7], $\mathsf{edge}(1, \mathsf{X}_0)$ is selected, so X_0 is substituted by the symbolic value c_0. A state transition labeled with $\mathsf{edge}(1, \mathsf{X}_0)$ is created, and a transition with the symbolic fact $\mathsf{edge}(1, \mathsf{c}_0)$ gives $\mathcal{S}_1$:

[9] $\mathsf{path}(1, \mathsf{c}_0)$. //Reduction of [7] with fact
[10] $\mathsf{answer}(\mathsf{c}_0)$. //Reduction of [6] with [9]
[11] $\mathsf{path}(1, \mathsf{X}_0) \leftarrow \mathsf{edge}(\mathsf{c}_0, \mathsf{X}_0)$. //Reduction of [8] with [9]
[12] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_0)$. //Copy of [6] because of [11]
[13] $\mathsf{path}(1, \mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_1) \wedge \mathsf{edge}(\mathsf{X}_1, \mathsf{X}_0)$. //Copy of [8] because of [11]

From $\mathcal{S}_1$ the successor state $\mathcal{S}_2$ is reached by a transition with the symbolic fact $\mathsf{edge}(\mathsf{c}_0, \mathsf{c}_1)$, labeled with the literal $\mathsf{edge}(\mathsf{c}_0, \mathsf{X}_0)$.

[14] $\mathsf{path}(1, \mathsf{c}_1)$. //Reduction of [11] with fact
[15] $\mathsf{answer}(\mathsf{c}_1)$. //Reduction of [12] with [14]
[16] $\mathsf{path}(1, \mathsf{X}_0) \leftarrow \mathsf{edge}(\mathsf{c}_1, \mathsf{X}_0)$. //Reduction of [13] with [14]
[17] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_0)$. //Copy of [12] because of [16]
[18] $\mathsf{path}(1, \mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_1) \wedge \mathsf{edge}(\mathsf{X}_1, \mathsf{X}_0)$. //Copy of [13] because of [16]

All following states resemble $\mathcal{S}_1$ but have symbolic constants $\mathsf{c}_2, \mathsf{c}_3, \ldots$ instead of c_1. □

For programs without recursive rules, the above state construction algorithm works well. In the other cases, as in our example, there will be infinitely many states since there are infinitely many symbolic values to be used at state transitions. Nevertheless, a part of these cases can be handled by trying to find finite many *equivalence classes* of states. It may be noticed that states are generated which have a similar structure but different symbolic values. Similar states generate again similar states because the same derivation steps are applied to similar sets of rules. Therefore they can be combined in one equivalence class of states.

Definition 15 (Equivalent States). *Let two states $\mathcal{S}_1$ and $\mathcal{S}_2$ be given. Let further $SYMB_1$ be the set of symbolic values occurring in $\mathcal{S}_1$ and $SYMB_2$ the set of symbolic values occurring in $\mathcal{S}_2$. $\mathcal{S}_1$ is equivalent to $\mathcal{S}_2$ iff a bijective mapping* map *from $SYMB_1$ to $SYMB_2$ exists so that the state $\mathcal{S}_1'$ obtained by replacing every symbolic value* v *in $\mathcal{S}_1$ by* $\mathsf{map}(\mathsf{v})$ *is equal to $\mathcal{S}_2$.*

The construction of an automaton with partial evaluation is straightforward. The states of this automaton represent equivalence classes of those states that are constructed during the derivation process. Consequently, when a state is constructed for which an equivalent state already exists, these states are fused to one state in the automaton.

Example 3. Consider again the transitive closure program with the goal rule and with symbolic values (Examples 1 and 2). It is visible that $\mathcal{S}_1$ and $\mathcal{S}_2$ are equivalent: if in $\mathcal{S}_2$ the symbolic value c_1 is mapped to c_0, both rule sets are equal. With this, the following state transition function results:

$\delta(\mathcal{S}_0\,,\, \mathsf{edge}(1, \mathsf{X}_0)) = \mathcal{S}_1$
$\delta(\mathcal{S}_1\,,\, \mathsf{edge}(\mathsf{c}_0, \mathsf{X}_0)) = \mathcal{S}_1$

□

In certain cases it happens that arbitrarily many different rules with the same structure of constants and variables but different values are accumulated in one state. In these cases the process of partial evaluation and automata construction does not terminate. A part of these cases results from tail recursive program rules. The problem with a tail recursive rule is that, starting from the last literal of the rule, arbitrarily long instantiation chains are created which have to be

kept in the state for reduction. These cases can be dealt with by introducing an additional derivation step and performing a *resolution* step instead of an instantiation when processing the last literal of a rule.

Definition 16 (Extension of Deduction Method). *The algorithm described in Definition 12 is extended as follows:*

1. *A reduction with a fact* F *the predicate of which is an EDB predicate is not affected and performed as described above.*
2. *Reductions with IDB facts are applied repeatedly, but only to rules* $\mathsf{R} = \mathsf{A} \leftarrow \mathsf{B}_1 \wedge \mathsf{B}_2 \wedge \cdots \wedge \mathsf{B}_n$ *where* $n > 1$.
3. *Instantiations are only applied to rules* $\mathsf{R} = \mathsf{A} \leftarrow \mathsf{B}_1 \wedge \cdots \wedge \mathsf{B}_n$ *with selected literal* B_1 *where the predicate of* B_1 *is an IDB predicate and* $n > 1$. *Otherwise, if* $n = 1$, *new rules are derived by applying a resolution step to* $\mathsf{A} \leftarrow \mathsf{B}_1$ *and program rules where the head literal unifies with* B_1. *The derivation step is therefore called* last literal resolution.
4. *If a rule* $\mathsf{R} = \mathsf{A} \leftarrow \mathsf{B}_1 \wedge \cdots \wedge \mathsf{B}_n$ *depends on a rule* R' *in the successor state it is only copied to the successor state if* $n > 1$.

The initial state consists of the goal rule and all rules that can be iteratively derived by instantiation and last literal resolution.

Example 4. Consider the tail recursive version of the transitive closure program:

[1] $\mathsf{path}(\mathsf{X}_0, \mathsf{X}_1) \leftarrow \mathsf{edge}(\mathsf{X}_0, \mathsf{X}_1)$.
[2] $\mathsf{path}(\mathsf{X}_0, \mathsf{X}_1) \leftarrow \mathsf{edge}(\mathsf{X}_0, \mathsf{X}_2) \wedge \mathsf{path}(\mathsf{X}_2, \mathsf{X}_1)$.

With the goal rule

[3] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_0)$.

the initial state is:

[4] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{path}(1, \mathsf{X}_0)$. //goal rule
[5] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{edge}(1, \mathsf{X}_0)$. //last literal resolution of [4]
[6] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{edge}(1, \mathsf{X}_1) \wedge \mathsf{path}(\mathsf{X}_1, \mathsf{X}_0)$. //last literal resolution of [4]

A transition with the symbolic fact $\mathsf{edge}(1, \mathsf{c}_0)$ gives $\mathcal{S}_1$:

[7] $\mathsf{answer}(\mathsf{c}_0)$. //reduction of [5] with fact
[8] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{path}(\mathsf{c}_0, \mathsf{X}_0)$. //reduction of [6] with fact
[9] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{edge}(\mathsf{c}_0, \mathsf{X}_0)$. //last literal resolution of [8]
[10] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{edge}(\mathsf{c}_0, \mathsf{X}_1) \wedge \mathsf{path}(\mathsf{X}_1, \mathsf{X}_0)$. //last literal resolution of [8]

The following states are equivalent to this state. The transition function is the same as for the left recursive program:

$\delta(\mathcal{S}_0\,,\, \mathsf{edge}(1, \mathsf{X}_0)) = \mathcal{S}_1$
$\delta(\mathcal{S}_1\,,\, \mathsf{edge}(\mathsf{c}_0, \mathsf{X}_0)) = \mathcal{S}_1$

□

In order to achieve termination, the other cases of infinitely growing states have to be excluded for the time being. For this the notion of the *schema* of a rule is introduced, and the set of *valid* states is restricted to those states that do not contain two rules with the same schema. PORTER uses a similar definition for a schema in [5].

Definition 17 (Schema of a Rule). *Let a normalized rule* R *be given, and let* $CONST_{\mathcal{V}}$ *be the set of those constants in* R *that do not occur in the program, which means they are symbolic values of the set* $\mathcal{V}$*. Let* $CONST_{\mathcal{V}}$ *be ordered by the occurrence of its elements in* R*:* $CONST_{\mathcal{V}} = \{c_i \mid i = 0, \ldots, |CONST_{\mathcal{V}}| - 1\}$ *where* $c_j < c_k$ *iff the first occurrence of* c_j *is before the first occurrence of* c_k *and* $j < k$ *iff* $c_j < c_k$ *(*$j, k \in \{0, \ldots, |CONST_{\mathcal{V}}| - 1\}$*). Let further* $\mathcal{B} = \{\mathsf{b}_i \mid i \in \mathbb{N} \cup \{0\}\}$ *be an ordered set of constants:* $\mathsf{b}_j < \mathsf{b}_k$ *iff* $j < k$ *(*$j, k \in \{\mathbb{N} \cup \{0\}$*). The schema of* R *is the rule obtained by replacing every* $c_i \in CONST_{\mathcal{V}}$ *by* $\mathsf{b}_i \in \mathcal{B}$ *(*$i \in \{0, \ldots, |CONST_{\mathcal{V}}| - 1\}$*).*

Example 5. Consider state $\mathcal{S}_1$ of the tail recursive program in Example 4. The schemata of the rules in this state are as follows:

[7] $\mathsf{answer}(\mathsf{b}_0)$.
[8] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{path}(\mathsf{b}_0, \mathsf{X}_0)$.
[9] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{edge}(\mathsf{b}_0, \mathsf{X}_0)$.
[10] $\mathsf{answer}(\mathsf{X}_0) \leftarrow \mathsf{edge}(\mathsf{b}_0, \mathsf{X}_1) \wedge \mathsf{path}(\mathsf{X}_1, \mathsf{X}_0)$.

□

Definition 18 (Valid State). *Let* $\mathcal{S}$ *be a state and* $\mathcal{SCH}$ *be the set of schemata of the rules in* $\mathcal{S}$*. The state* $\mathcal{S}$ *is valid iff* $|\mathcal{SCH}| = |\mathcal{S}|$*.*

For a given program that meets all requirements mentioned above there are only finitely many possibilities to create valid states, so partial evaluation is guaranteed to terminate.

An implementation of the automaton will use states where the symbolic values are replaced by assignable variables that hold the actual data values. Explicit constants can be included in the target code and no derivations have to be performed so that the runtime states are very compact and the main task at state transitions should be accessing and selecting the data. Different results for the answer predicate can be obtained by backtracking or by concurrent processing of alternative transitions.

5 Conclusion

We have presented an algorithm for efficient query evaluation and preprocessing of function-free logic programs based on Earley Deduction. The algorithm can process non-recursive, left- and tail-recursive rules and has been proven to be correct, complete, and terminating. While Earley Deduction can in principle be used for arbitrary logic programs, still the basic algorithm presented here is already an improvement to it because it processes several derivations for one

fact in one step. There is a special optimization potential when it can be proven that only one fact is applicable in a state and we do not have to check whether there is a cycle in the state sequence. Subjects of our future research include further optimizations for special applications and an efficient implementation of the generated automaton.

Further material, including a demo program (written in SWI Prolog) showing the state sequences for a given program and query, is available at `http://dbs.informatik.uni-halle.de/Earley`.

References

1. Brass, S.: SLDMagic — the real magic (with applications to web queries). In: Lloyd, W., et al. (eds.) First International Conference on Computational Logic (CL'2000/DOOD'2000). pp. 1063–1077. No. 1861 in LNCS, Springer, Heidelberg, Berlin (2000)
2. Brass, S.: Implementation alternatives for bottom-up evaluation. In: Hermenegildo, M., Schaub, T. (eds.) Technical Communications of the 26th International Conference on Logic Programming (ICLP'10). Leibniz International Proceedings in Informatics (LIPIcs), vol. 7, pp. 44–53. Schloss Dagstuhl (2010), http://drops.dagstuhl.de/opus/volltexte/2010/2582
3. Brass, S.: Order in datalog with applications to declarative output. In: Barcelo, P., Pichler, R. (eds.) Datalog 2.0. LNCS, vol. 7494, pp. 56–67. Springer-Verlag (2012)
4. Pereira, F.C.N., Warren, D.H.D.: Parsing As Deduction. In: Proceedings of the 21st annual meeting on Association for Computational Linguistics. pp. 137–144. ACL '83, Association for Computational Linguistics, Stroudsburg, PA, USA (1983), `http://dx.doi.org/10.3115/981311.981338`, http://www.aclweb.org/anthology/P83-1021
5. Porter III, H.H.: Optimizations to Earley Deduction for DATALOG Programs (1985), http://web.cecs.pdx.edu/~harry/earley/datalog.pdf
6. Porter III, H.H.: Earley Deduction (1986), http://web.cecs.pdx.edu/~harry/earley/earley.pdf